arXiv:1704.04097v1 [cs.CV] 13 Apr 2017

Recognizing Activities of Daily Living from Egocentric Images

Alejandro Cartas[1⋆], Juan Marín[1⋆], Petia Radeva[1,2], and Mariella Dimiccoli[1,2]

[1] University of Barcelona, MAIA Department, 08007 Barcelona, Spain,
E-mail: alejandro.cartas@ub.edu, jmarinve7@alumnes.ub.edu
[2] Computer Vision Center, 08193 Cerdanyola del Vallès, Barcelona, Spain

Abstract. Recognizing Activities of Daily Living (ADLs) has a large number of health applications, such as characterize lifestyle for habit improvement, nursing and rehabilitation services. Wearable cameras can daily gather large amounts of image data that provide rich visual information about ADLs than using other wearable sensors. In this paper, we explore the classification of ADLs from images captured by low temporal resolution wearable camera (2fpm) by using a Convolutional Neural Networks (CNN) approach. We show that the classification accuracy of a CNN largely improves when its output is combined, through a random decision forest, with contextual information from a fully connected layer. The proposed method was tested on a subset of the NTCIR-12 egocentric dataset, consisting of 18,674 images and achieved an overall accuracy of 86% activity recognition on 21 classes.

Keywords: egocentric vision, lifelogging, activity recognition, convolutional neural networks

1 Introduction

The Activities of Daily Living (ADLs) include, but are not limited to the activities that an independent person performs on daily basis for living at home or in a community [8]. The monitoring of these activities on elderly people could prevent health problems [8,11]. Recently, egocentric (first-person) cameras have been used to monitor ADLs, because they can provide richer contextual information than using only traditional sensors [9]. These lifelogging devices can generate large volumes of data in matter of days. For instance, a wearable camera such as the Narrative Clip can take more than 2,800 pictures per day from an egocentric (first-person) point of view. To extract information about the behavior of a user, these data needs to be classified in an orderly and timely fashion.

Over the last five years, egocentric activity recognition has been an active area of research. Fathi et al. [2] presented a probabilistic model that classifies activities from short egocentric videos. Their approach models an activity into

⋆ Both authors contributed equally.

Fig. 1: Examples of annotated images with their corresponding activity label.

a set of different actions, and each action is modeled per frame as a spatio-temporal relationship between the hands and the objects involved on it. They further extend their work [3] by another probabilistic generative model that incorporates the gaze features and that models an action as a sequence of frames. Pirsiavash and Ramanan [10] introduced a dataset of 18 egocentric actions of daily activities performed by 20 persons in unscripted videos. Furthermore, they presented a temporal pyramid to encode spatio-temporal features along with detected active objects knowledge. These temporal pyramids are the input of support vector machines trained for action recognition. More recently, Ma et al. [7] proposed a twin stream Convolutional Neural Network (CNN) architecture for activity recognition from videos. One of the streams is used for recognizing the appearance of an object based on a hand segmentation and a region of interest. The other stream recognizes the action using an optical flow sequence. In order to recognize activities, both streams are join and the last layers are finetuned.

Activity classification from egocentric photo streams is even more difficult problem than from video, since they provide less contextual action information. Castro et. al. [1] introduced a lifelogging dataset composed of 40,103 egocentric images from one subject taken during a 6 months period. The images from this dataset were annotated in 19 different activity categories. In this dataset, the activities performed by the user tend to be routinary and in the same environment. In other words, the activities are almost performed daily at the same time and involved the same objects. Consequently, time and global image features such as color convey useful information for describing activities. To exploit both characteristics, the authors introduced the Late Fusion Ensemble (LTE) method that combines, through a random decision forest, the classification probabilities of a CNN with time and global features, namely color histogram, from the input image. The shortcoming of this approach is that it cannot be generalized to multiple users, since the network need to learn contextual information for each

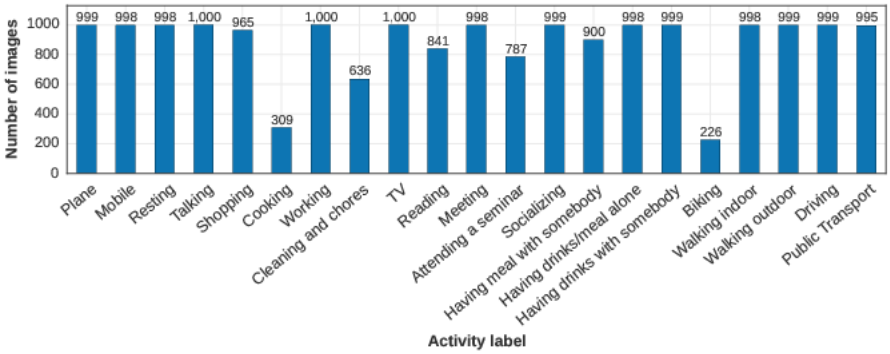

Fig. 2: Our dataset summary.

new user. For instance, two distinct persons might have different daily routines, depending on their job, age, etc.

In this work, we make a step forward on activity recognition from egocentric images by generalizing the task of activity recognition to multiple users. Our approach combines the outputs of different layers of a CNN and use them as the input of a random decision forest. We tested our method on a subset of the NTCIR-12 egocentric dataset [4], consisting of 18,674 pictures acquired by three users and annotated with 21 activity labels. Some examples of labeled images with their corresponding category are shown in Fig. 1. Our approach is similar to the LTE introduced in [1], in the sense that it uses a random forest to combine the output of a CNN with contextual information. However, there are several differences that should be highlighted. First, the problem itself we face is different: instead of classifying the activities of a single user, we classify the activities of threee different users from a dataset having less than half the number of images of the one used in [1]. Hence, our task is more challenging because we must deal with an increased intra-class variability with a much smaller number of images. In addition, since the pictures from the dataset were taken by users with different lifestyles, our method cannot take advantage of their time information and color histogram. To face these problems, instead of using time and color as contextual information, we use the output of a fully connected layer of the CNN.

The rest of the paper is organized as follows: in section 2 we provide an overview of our approach. In section 3, we introduce the dataset used in the experiments and we briefly describe the annotation tool we created for labeling. We detail the methodology we followed for conducting our experiments and the different combinations of networks and layers we used in section 4. The results we obtained are discussed in section 5. Finally, in section 6, we present our conclusion and final remarks.

2 Activity classification

Our activity classification method is an ensemble composed of a CNN and a random forest. Specifically, our approach joins the output vectors from two layers of a CNN, i.e. the fully connected and the softmax layers, and gives them as input

to a random forest. The training of the ensemble is a two-step process. First, a CNN is finetuned and then a random forest is trained over its output vectors.

The objective of our experiments was to show how the classification performance of the CNN improves by using contextual information. In these experiments, we used two networks as the base of our ensembles, namely the AlexNet [6] and GoogLeNet [12]. Both neural networks were finetuned on a subset of images from the NTCIR-12 dataset [4] annotated for activity classification.

3 Dataset

In our experiments, we used a subset of images from the NTCIR-12 dataset [4] that consists of 89,593 egocentric pictures belonging to three persons. Each user worn a first-person camera in a period of almost a month, totaling 79 days. This camera passively took two pictures per minute.

The dataset we used in this work, consists of a subset of 18,674 images from the NTCIR-12 dataset. These pictures correspond to all users and have different dates and times. We used 21 activity categories to label them using our annotation tool, which is briefly detailed in the next subsection. Although our subset of annotated images is imbalanced, less than the half of categories have less instances than the rest. Some examples of annotated activities and the distribution of the number of images by category are shown in Fig. 1 and Fig. 2, respectively. The dataset was split in 13,991 training images, 1,857 validation images, and 2,796 testing images. This split maintained the same proportion of examples per category and was used in all our experiments.

Annotation tool We created a web-based annotation tool specifically target for labeling large amounts of images. The following design guidelines were considered:

- *Easy interaction.* Browsing and annotating large number of images can be done in an intuitive way. Additionally, descriptive tags can be individually added to a specific picture.
- *Speed and performance.* The tool can handle several connected users at the same time and present their large collection of pictures.
- *Privacy and security.* Personal pictures of a user are maintained private.

4 Methodology

The main objective of the experiments was to determine if contextual information can improve the activity classification accuracy of a CNN by ensembling the CNN with a random forest. Additionally, we tried to find out the ensemble with the best combination of layers for performance improvement. For training and testing purposes we used the dataset split described in the last section. Since

Activity	AlexNet	AlexNet+RF on FC6	AlexNet+RF on Prob	AlexNet+RF on FC6+Prob	GoogLeNet	GoogLeNet+RF on Prob	GoogLeNet+RF on pool5/7x7+prob
Public Transport	82.17	**87.60**	86.05	86.82	79.07	80.62	84.50
Driving	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**	**100.00**
Walking outdoor	84.52	86.31	**88.10**	86.31	83.93	85.12	**88.10**
Walking indoor	61.88	68.75	65.00	**71.88**	60.00	63.75	68.75
Biking	**81.58**	73.68	78.95	73.68	68.42	71.05	76.32
Drinking together	84.09	**90.91**	82.58	87.88	76.52	80.30	87.88
Drinking/eating alone	70.89	**84.81**	74.05	80.38	60.13	70.25	75.95
Eating together	76.56	**86.72**	80.47	82.81	75.78	82.03	85.16
Socializing	65.31	**89.80**	77.55	83.67	63.95	72.11	79.59
Attending a seminar	89.92	**93.28**	89.08	**93.28**	82.35	86.55	89.08
Meeting	78.01	85.82	80.14	85.11	73.76	82.98	**87.94**
Reading	69.40	84.33	77.61	82.09	85.82	83.58	**86.57**
TV	88.74	95.36	92.72	94.04	94.70	96.03	**98.01**
Cleaning and chores	44.74	61.40	56.14	71.93	51.75	67.54	**77.19**
Working	89.24	**95.57**	91.14	**95.57**	90.51	91.77	94.94
Cooking	31.11	46.67	**48.89**	42.22	24.44	37.78	44.44
Shopping	73.88	**85.07**	76.12	81.34	76.87	76.87	78.36
Talking	70.47	82.55	77.85	81.88	71.81	78.52	**83.89**
Resting	98.70	**99.35**	98.70	**99.35**	97.40	98.70	**99.35**
Mobile	79.19	85.91	85.23	**88.59**	79.87	84.56	86.58
Plane	90.91	93.01	91.61	90.91	93.71	94.41	**96.50**

Accuracy	78.51	**86.58**	82.36	85.79	78.11	82.22	86.04
Macro precision	78.46	**87.29**	81.57	86.83	77.17	81.55	86.20
Macro recall	76.73	**84.61**	80.86	83.80	75.75	80.22	84.24
Macro F1-score	77.09	**85.45**	81.11	84.68	76.06	80.64	84.84

Table 1: Comparison of the ensembles of CNN+Random forest on different combinations of layers. Upper table shows the recall per class and the lower table shows the performance metrics.

the dataset was imbalanced, we assessed the performance of the ensembles by not only using the accuracy, but also macro metrics for precision, recall, and F1-score. The next subsection presents all the proposed ensemble configurations with their respective training procedure.

4.1 Ensemble configurations

In order to train our ensembles, we first fine-tuned AlexNet and GoogLeNet on our dataset. Both models were fine-tuned using the Caffe framework [5] with the same number of iterations (approx. 10 epochs) and the following settings:

- **AlexNet**. It was trained using stochastic gradient descent for 2,180 iterations with a batch size of 64 images, a learning rate $\alpha = 3\text{x}10^{-5}$, and a momentum $\mu = 0.9$.
- **GoogLeNet**. It was trained for 4,370 iterations with a batch size of 32 images, a learning rate $\alpha = 3.7\text{x}10^{-5}$, and a momentum $\mu = 0.9$.

After fine-tuning our baseline networks, we train several ensembles of random forests combining different final output layers of the CNNs. The following combinations were used in our experiments:

- **AlexNet+RF on Prob**. This configuration was trained using a random forest of 500 trees on the softmax probability layer.
- **GoogLeNet+RF on Prob**. This configuration was trained using a random forest of 500 trees on the softmax probability layer.
- **AlexNet+RF on FC6**. We trained a random forest of 500 trees on only the first fully connected layer (FC6), containing a vector of 4,096 elements.

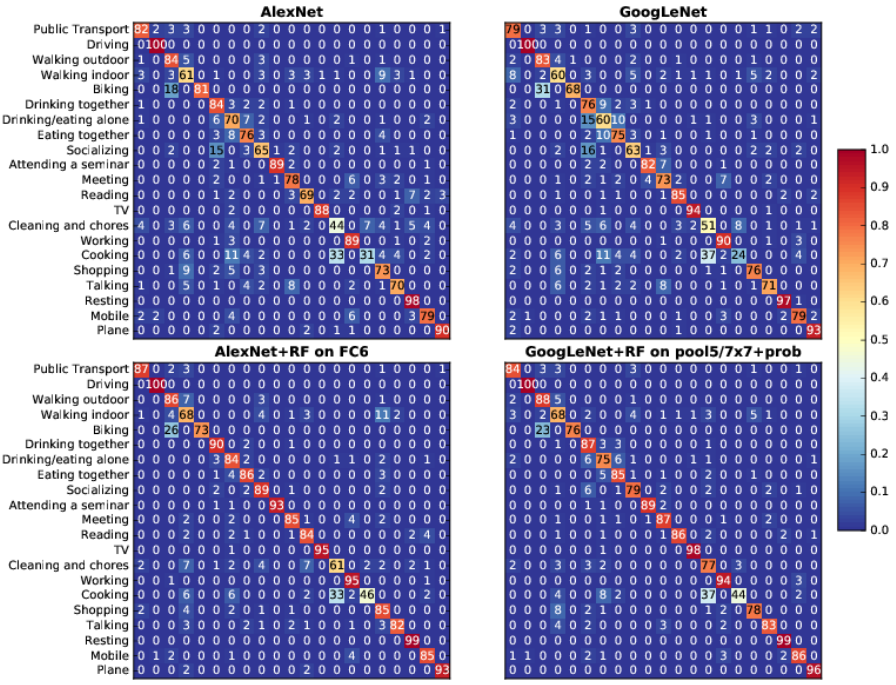

Fig. 3: Normalized confusion matrices of the best combination of layers for each baseline convolutional neural network. This figure is best seen in color.

- **AlexNet+RF on FC6+Prob**. A random forest of 500 trees was trained on the FC6 and the softmax layers of AlexNet, thus summing a vector of 4,117 elements.
- **GoogLeNet+RF on Pool5/7x7+Prob**. This configuration was trained using a random forest of 500 trees on the pool5/7x7 fully connected layer and the softmax probability layer, thus having a vector of 1,045 features.

5 Results

Our experiments show that the ensemble of CNN plus a random forest improves the performance for both baseline CNN. Table 1 shows the classification performance on the the baseline CNN and on different ensembles. The baseline CNN have a similar performance. After using the random forest on the output the softmax layer, the accuracy improved around 4%. Specifically, the recall of some categories with fewer learning instances improved significantly, such as *Cooking*, and *Cleaning and Choring*. Additionally, high overlapping classes such as *Drinking/Eating alone* and *Eating together* also improved their accuracy. The improvement over the overlapping classes can also be seen the confusion matrices shown in Fig. 3. This means that the random forest improved the classification

Drinking/eating alone

AlexNet Top 5

#	Activity	Score
1	Public Transport	0.1864
2	Cooking	0.1464
3	Eating together	0.1382
4	**Drinking/eating alone**	**0.1223**
5	Cleaning and chores	0.1067

GoogLeNet Top 5

#	Activity	Score
1	Plane	0.2004
2	Public Transport	0.1943
3	Cleaning and chores	0.1450
4	Cooking	0.0925
5	**Drinking/eating alone**	**0.0748**

Meeting

AlexNet Top 5

#	Activity	Score
1	Talking	0.2034
2	**Meeting**	**0.1701**
3	Cooking	0.1090
4	Cleaning and chores	0.0903
5	Shopping	0.0878

GoogLeNet Top 5

#	Activity	Score
1	**Eating together**	**0.3065**
2	Talking	0.1215
3	Socializing	0.1062
4	Cleaning and chores	0.08126
5	Meeting	0.0664

Reading

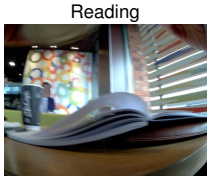

AlexNet Top 5

#	Activity	Score
1	**Drinking/eating alone**	**0.5002**
2	Cleaning and chores	0.1511748880
3	Eating together	0.1263086796
4	Shopping	0.0589886233
5	Drinking together	0.0251834411

GoogLeNet Top 5

#	Activity	Score
1	**Cleaning and chores**	**0.4259**
2	Eating together	0.1145323068
3	Drinking/eating alone	0.1137253270
4	Drinking together	0.0841688886
5	Reading	0.06203

Fig. 4: Classification activity examples. On top of each image is shown its true activity label and on bottom its top 5 predictions by AlexNet and GoogLeNet. Additionally, the result of the ensembles *AlexNet+RF on FC6* and *GoogLeNet+RF on Pool5/7x7+prob* is highlighted on color in its corresponding table. The green and red colors means true positive and false positive classification, respectively.

of images belonging to categories that score similar probabilities. Moreover, the only decrease on accuracy is presented on the class *Biking*. Since its accuracy on the baseline CNN is very high (81.68%) considering the small number of learning instances (226), we believe this decrease is a consequence of the random forest trying to balance the prediction error among classes.

The results of adding contextual information from fully connected layers show a better performance on classification. Table 1 shows that the best ensembles were the *AlexNet+RF on FC6* and *GoogLeNet+RF on Pool5/7x7+prob*. Furthermore, these ensembles improved the baseline accuracy by 8.07% and 7.93% for AlexNet and GoogLeNet, respectively. Although the performance metrics improved decreasingly with respect to the ensembles that only used the softmax layer, the extra features removed the overfitting problem on the *Walking outdoor* category. Some classification examples are shown in Fig. 4.

6 Conclusion

We presented an egocentric activity classifier ensemble method, which combines different layers of a CNN through a random forest. Specifically, the random forest takes as input a vector containing the output of the softmax probability layer and a fully connected layer encoding global image features. We tested several ensembles based on AlexNet and GoogLeNet, achieving a 8% performance improvement on our dataset. The proposed method has been tested on a subset of the NTCIR-12 egocentric dataset consisting of 18,674 images that we

labeled with 21 different activity labels. Although we obtained 86% accuracy on a quite difficult task, we believe that there is still room for improvement. Indeed, the proposed approach operates at image-level, without taking into account the temporal coherence of photo-streams. Future work will investigate how to take temporal coherence into account.

Acknowledgments

A.C. was supported by a doctoral fellowship from the Mexican Council of Science and Technology (CONACYT) (grant-no. 366596). This work was partially founded by TIN2015-66951-C2, SGR 1219, CERCA, *ICREA Academia'2014* and 20141510 (Marató TV3). The funders had no role in the study design, data collection, analysis, and preparation of the manuscript. M.D. is grateful to the NVIDIA donation program for its support with GPU card.

References

1. D. Castro, S. Hickson, V. Bettadapura, E. Thomaz, G. Abowd, H. Christensen, and I. Essa. Predicting daily activities from egocentric images using deep learning. In *proceedings of the 2015 ACM International symposium on Wearable Computers*, pages 75–82. ACM, 2015.
2. A. Fathi, A. Farhadi, and J. M. Rehg. Understanding egocentric activities. In *2011 International Conference on Computer Vision*, pages 407–414. IEEE, 2011.
3. A. Fathi, Y. Li, and J. M. Rehg. Learning to recognize daily actions using gaze. In *European Conference on Computer Vision*, pages 314–327. Springer, 2012.
4. C. Gurrin, H. Joho, F. Hopfgartner, L. Zhou, and R. Albatal. Overview of ntcir-12 lifelog task. In *NTCIR-12*. National Institute of Informatics (NII).
5. Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
6. A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
7. M. Ma, H. Fan, and K. M. Kitani. Going deeper into first-person activity recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
8. I. Martin-Lesende, K. Vrotsou, I. Vergara, A. Bueno, A. Diez, et al. Design and validation of the vida questionnaire, for assessing instrumental activities of daily living in elderly people. *J Gerontol Geriat Res*, 4(214):2, 2015.
9. T.-H.-C. Nguyen, J.-C. Nebel, and F. Florez-Revuelta. Recognition of activities of daily living with egocentric vision: A review. *Sensors (Basel)*, 16(1):72, Jan 2016. sensors-16-00072[PII].
10. H. Pirsiavash and D. Ramanan. Detecting activities of daily living in first-person camera views. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2847–2854. IEEE, 2012.

11. S. M. Schüssler-Fiorenza Rose, M. G. Stineman, Q. Pan, H. Bogner, J. E. Kurichi, J. E. Streim, and D. Xie. Potentially avoidable hospitalizations among people at different activity of daily living limitation stages. *Health services research*, 2016.
12. C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015.